REMARKS

This Amendment is filed in response to the FINAL Office Action mailed July 19, 2010. A Request for Continued Examination and the associated fee is also filed herewith. All objections and rejections are respectfully traversed.

Claims 1-8, 12, 14-25, 29, 31-35, 37-39, 41-42, 46-50 are in the case.

No new claims have been added.

Claims 1-8, 12, 14-22, 24-25, 29, 31-32, 34-35, 38-39, 42, and 46-50 have been amended.

Interview Summary

Applicant would like to thank Examiner Lovel for conducting the Applicant Initiated Interview on October 7, 2010 and for helping to advance this Application closer to allowance. Generally, as will be elaborated upon in greater detail below, the issue discussed involved Applicant's claimed novel and non-obvious **management application to scan the directory tree *through at least one volume information block configured to reference each root of each PCPI comprising the directory tree* to generate an index associated with the directory tree**. Examiner Lovel indicated that the claims as discussed overcame the prior art of record. Should this Amendment not place all claims in condition for allowance, Applicant respectfully requests that Examiner contact the undersigned attorney to discuss any remaining issues *before* issuing a next Office Action. Examiner is encouraged to contact the undersigned attorney with any questions.

Claim Objections

At paragraph 5 of the Office Action, Examiner objected to claims 15 and 32 were objected to due to informalities. Applicant thanks Examiner for noting such informalities. Applicant acknowledges Examiner's recommendations to change the dependencies of claims 15 and 32. However, Applicant respectfully notes that claims 15 and 32 do not require the amendment suggested by Examiner.

Rejections Under 35 U.S.C. §103

At paragraph 10 of the Office Action, claims 1-2, 6, 15-19, 31-35, 38, and 49-50 were rejected under 35 U.S.C. §103(a) as being obvious over Mani-Meitav et al., U.S. Patent Application Publication No. 2005/0216788 published on September 29, 2005 (hereinafter "Mani"), in view of Dewey et al., U.S. Patent No. 7,529,778 issued on May 5, 2009 (hereinafter "Dewey").

Applicant's claimed novel and non-obvious invention, as set forth in representative claim 1, comprises in part:

> The present invention, as set forth in representative claim 1, comprises in part:
>
> 1. A system for indexing and manipulating backup data stored on a destination storage system, comprising:
>
> one or more source storage systems configured to transmit the backup data to the destination storage system;
>
> a management client comprising a processor configured to execute a management application, the management application configured to communicate with the destination storage system and further configured to access data identifiers related to the backup data organized in a directory tree representing a plurality of persistent consistency point images (PCPIs) of the backup data, wherein each PCPI is associated with a creation time, the **management application** further configured **to scan the directory tree *through at least one volume information block configured to reference each root of each PCPI comprising the directory tree* to generate an index** of directories, files, or qtrees **associated with the directory tree**, the management application further configured to organize the data identifiers to enable the backup data to be displayed on a display screen of the management client; and
>
> a user interface of the management client configured to select a directory, file, or qtree to view, and, in response to the selection, the management client further configured to query the management application and in response to the query return a list of the selected directory, file, or qtree and one or more versions of the selected directory, file, or qtree.

Mani teaches cataloguing snapshots so that in recovery mode, "a chosen version of a certain set of data may be selected" [0138]. Notably, Mani teaches assembling an

image of a desired file by analyzing the data blocks residing in a repository. Specifically, a "file-level recovery procedure...analyze[s] block-level data for deriving therefrom both file-system structure and file-system organization. The analysis process scans the raw block-level backup image in the repository for file-system organization" [0166].

Dewey teaches obtaining a list of earlier file versions that exist on a temporal shadow volume (col. 8, lines 60-65). Specifically, to obtain such a list, Dewey explicitly states that locating the file versions is accomplished by first obtaining at the local (client) machine a list of the available shadow volumes (col. 9, lines 23-28; *see also* Fig. 3). If any shadow volume exists, this information is returned in the shadow volume list (col. 9, lines 49-51).

Applicant respectfully urges that Mani, taken singly or in any combination with Dewey, does not disclose, teach, or suggest Applicant's claimed novel and non-obvious **management application to scan the directory tree *through at least one volume information block configured to reference each root of each PCPI comprising the directory tree* to generate an index associated with the directory tree**.

Applicant claims, in part, a **management application** configured to communicate with a destination storage system and further configured to access data identifiers related to the backup data organized in a directory tree structure representing a plurality of persistent consistency point images (PCPIs) of the backup data. Each PCPI may pertain to a particular directory tree or, e.g., file. It may be desirable to generate an index of those PCPIs associated with a particular directory tree spanning multiple points-in-time (i.e., ***each PCPI comprising the directory tree***), rather than generating an index of every PCPI making up a single point-in-time image of the entire active file system. In other words, where information on a particular directory is requested, the full range of those PCPIs containing that particular directory (e.g., the root of each relevant PCPI comprising that particular directory tree) may be located and read to generate an index associated with

that particular directory tree. Applicant illustratively accomplishes this, in part, by using a **management application to scan the directory tree** ***through at least one volume information block configured to reference each root of each PCPI comprising the*** (particular) ***directory tree*** **to generate an index associated with the** (that particular) **directory tree**.

Applicant respectfully contends that Mani fails to teach or suggest Applicant's claimed novel and non-obvious **management application to scan the directory tree** ***through at least one volume information block configured to reference each root of each PCPI comprising the directory tree*** **to generate an index associated with the directory tree**. Specifically, while Mani teaches cataloguing "snapshots", Mani teaches that the cataloguing is generated by reading and retrieving data stored in a repository in the block-level data configuration and scanning the raw block-level backup image in the repository by reading the type of file-system from the image. Notably, neither of those steps is taught or suggested to be the same as scanning **the directory tree** ***through at least one volume information block***, nor does Mani teach or suggest that Mani's raw block-level backup image is the same as the ***at least one volume information block configured to reference each root of each PCPI comprising the directory tree***. Moreover, Mani is silent on the concept of a volume information block. In contrast, Applicant claims a management application configured to scan the directory tree ***through at least one volume information block configured to reference each root of each PCPI comprising the directory tree*** to generate an index associated with the directory tree. As such, because Mani catalogues "snapshots" by methods other than scanning ***at least one volume information block configured to reference each root of each PCPI comprising the directory tree***, Mani fails to teach or suggest Applicant's claimed novel and non-obvious **management application to scan the directory tree** ***through at least one volume information block configured to reference each root of each PCPI comprising the directory tree*** **to generate an index associated with the directory tree**.

Applicant respectfully contends that Dewey fails to teach or suggest Applicant's claimed novel and non-obvious **management application to scan the directory tree *through at least one volume information block configured to reference each root of each PCPI comprising the directory tree* to generate an index associated with the directory tree**. Specifically, while Dewey teaches obtaining a list of earlier file versions, Dewey also states that such a list is generated by obtaining at the local (client) machine a list of all the available shadow volumes. In other words, if any shadow volume exists, this information is returned in the shadow volume list. However, none of Dewey's disclosed details of how the file versions are located are taught or suggested to be the same as Applicant's claimed scanning ***at least one volume information block configured to reference each root of each PCPI comprising the directory tree***. In contrast, Applicant claims a management application configured to scan the directory tree ***through at least one volume information block configured to reference each root of each PCPI comprising the directory tree*** to generate an index associated with the directory tree. As such, because Dewey teaches generating a list of files by methods other than scanning ***at least one volume information block configured to reference each root of each PCPI comprising the*** (requested) ***directory tree***, Dewey fails to teach or suggest Applicant's claimed novel and non-obvious **management application to scan the directory tree *through at least one volume information block configured to reference each root of each PCPI comprising the directory tree* to generate an index associated with the directory tree**.

Based on the foregoing, Applicant respectfully urges that Mani, taken singly or in any combination with Dewey, fails to render the presently claimed invention obvious under 35 U.S.C. §103(a). Mani and Dewey, taken singly or in any combination, do not disclose, teach, or suggest Applicant's claimed novel and non-obvious **management application to scan the directory tree *through at least one volume information block configured to reference each root of each PCPI comprising the directory tree* to generate an index associated with the directory tree**.

At paragraph 11 of the Office Action, claims 3-5, 20-23, 37, and 41 were rejected under 35 U.S.C. §103(a) as being obvious over Mani, in view of Dewey, and in further view of Armangau, U.S. Patent No. 6,434,681 (hereinafter "Armangau").

Applicant respectfully notes that claims 3-5, 20-23, 37, and 41 are dependent claims that depend from independent claims believed to be in condition for allowance. Accordingly, claims 3-5, 20-23, 37, and 41 are believed to be in condition for allowance.

At paragraph 12 of the Office Action, claims 7-8, 12, 14, 24-25, and 29 were rejected under 35 U.S.C. §103(a) as being obvious over Mani, in view of Dewey, and in further view of Arakawa et al., U.S. Patent Application Publication No. 2003/0131207 (hereinafter "Arakawa").

Applicant respectfully notes that claims 7-8, 12, 14, 24-25, and 29 are dependent claims that depend from independent claims believed to be in condition for allowance. Accordingly, claims 7-8, 12, 14, 24-25, and 29 are believed to be in condition for allowance.

At paragraph 13 of the Office Action, claims 42 and 46-48 were rejected under 35 U.S.C. §103(a) as being obvious over Mani, in view of Dewey, and in further view of Arakawa.

The present invention, as set forth in representative claim 42, comprises in part:

42. A system, comprising:

a source storage system configured to generate a plurality of persistent consistency point images (PCPIs) associated with a particular directory tree, and further configured to transfer the plurality of PCPIs to a destination storage system;

the destination storage system configured to execute a management client, the management client comprising a processor configured to

execute a management application, the **management application** configured **to scan the particular directory tree** ***through at least one volume information block configured to reference each root of each PCPI comprising the particular directory tree*** **to organize the plurality of PCPIs into an index** using a database operatively connected to the management client configured to allow the plurality of PCPIs to be displayed on a display screen of the management client as a listing of source data entries indexed by the particular directory tree, each PCPI of the particular directory tree created at one or more different creation times, and to allow the plurality of PCPIs to be displayed on the display screen as a listing of source data entries indexed by names of the source storage system, and to allow the plurality of PCPIs to be displayed on the display screen as a listing of source data entries indexed by names of volumes of the destination storage system in which backup data from the source storage system resides; and

an interface of the management client configured to select a data entry for the particular directory tree, and, in response to the selection, query the management application and in response to the query return a list of the plurality of PCPIs associated with the particular directory tree.

Arakawa teaches a virtualized volume snapshot formation method. However, Arakawa fails to teach or suggest Applicant's claimed novel and non-obvious **management application** configured **to scan the particular directory tree** ***through at least one volume information block configured to reference each root of each PCPI comprising the particular directory tree*** **to organize the plurality of PCPIs into an index**.

Additionally, as noted above with regard to claim 1, Mani and/or Dewey also do not teach or suggest Applicant's claimed novel and non-obvious **management application to scan the directory tree** ***through at least one volume information block configured to reference each root of each PCPI comprising the directory tree*** **to generate an index associated with the directory tree**. As such, *inter alia,* because claim 42 comprises similar limitations of claim 1 not shown or suggested by any prior art reference (or any combination thereof), Applicant respectfully contends that Mani, taken singly or in any combination with Dewey and/or Arakawa, fails to render the presently claimed invention obvious under 35 U.S.C. §103(a). Specifically, Mani and/or Dewey and/or Ara-

kawa, taken singly or in any combination, do not teach or suggest Applicant's claimed novel and non-obvious **management application** configured **to scan the particular directory tree *through at least one volume information block configured to reference each root of each PCPI comprising the particular directory tree* to organize the plurality of PCPIs into an index**.

Applicant's Interpretation of the Prior Art

Applicant's interpretation of the prior art was derived, in part, from the following excerpts:

Mani

[0166] The FBSRD operates a file-level recovery procedure enabling the restoration of data to file-level data configuration, by reading and retrieving data stored in the repository 15 in the block-level data configuration…The analysis process scans the raw block-level backup image in the repository 15 for file-system organization, by first reading the type of file-system from the image, and then proceeds by applying a specific analysis method for each type of file-system. The analysis process proceeds iteratively through all elements of the file-system structure until all files and directories have been analyzed. The derived file-system structure is stored in the repository 15 as a catalog database for subsequent use during file-level restoration, when in recovery mode. Each catalog is associated with a specific snapshot, and is used to find data related to one snapshot. (emphasis added)

Dewey

…locating the file versions is accomplished by first obtaining at the local (client) machine a list 300 of the available shadow volumes. To this end, the version locator API 208 includes or is otherwise associated with a client-side shadow volume location mechanism 304, which communicates with a counterpart server-side volume location mechanism 320 to obtain the list 300 of the temporal shadow volumes. (emphasis added)

Then, the shadow volume manager 322 is asked to return information (e.g. a timestamp) for the identified base volume (e.g., 228) for each shadow volume that has been captured therefor. If any shadow volumes exist, this information is then returned in the shadow volume list 300. (col. 9, lines 49-51) (emphasis added)

Conclusion

All claims are believed to be fully supported by Applicant's specification.

All independent claims are believed to be in condition for allowance.

All dependent claims are believed to be dependent from allowable independent claims, and therefore in condition for allowance.

Favorable action is respectfully solicited.

Please charge any additional fee occasioned by this paper to our Deposit Account No. 03-1237.

Respectfully submitted,

/Michael T. Abramson/
Michael T. Abramson
Reg. No. 60,320
CESARI AND MCKENNA, LLP
88 BLACK FALCON AVENUE
BOSTON, MA 02210
Telephone: (617) 951-2500
Facsimile: (617) 951-3927